Partnership for Multiscale Gyrokinetic (MGK) Turbulence – FASTMath and RAPIDS Applications

D. R. Ernst (MIT)

FASTMath All-Hands Meeting, June 10-11 (2019), Argonne

SciDAC:

David Hatch [Principal Investigator], M. Kotschenreuther, F. Jenko, C. Michoski (U. Texas-Austin); L. LoDestro, J. Parker (LLNL); W. Dorland (U. Maryland); G. Hammett, A. Hakim (PPPL); D. Ernst (MIT)

FASTMath:

D. Reynolds (SMU)

C. Woodward, D. Gardner (LLNL)

M. Minion (LBNL)

RAPIDS:

S. Hongzhang, Lenny Oliker (LBNL)

Courant:

A. Cerfon (NYU)

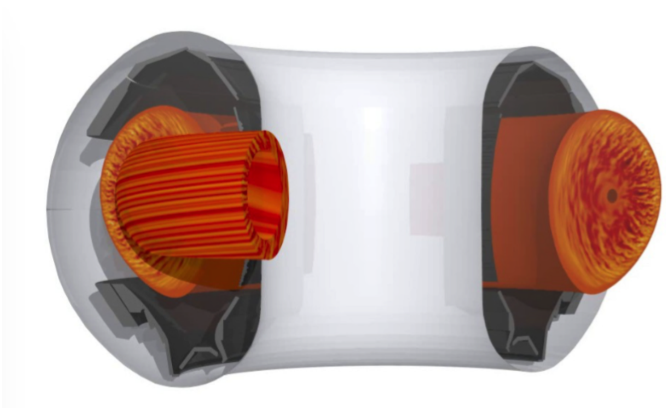

SciDAC MGK

mgkscidac.org

MGK Partnership Focused on Multiscale Challenges

Simulations of *core* turbulence in Tokamak Fusion Plasmas have matured:

- **Simulations have matched Particle, Momentum, Energy fluxes and Turbulence Fluctuation Spectra *simultaneously* in experiments, within error bars**
  [e.g., Ernst et al. Phys. Plasmas (2016)]
- **Turbulent transport limits fusion power $\propto$ (pressure)^2**

Multiscale challenges:

- **Turbulence in the sharp edge controls the core, but hard to simulate**
- **Electron thermal transport can involve turbulence scales spanning two orders of magnitude**
- **Evolving turbulence on 100x longer transport timescale to find temperature, density, rotation profiles**

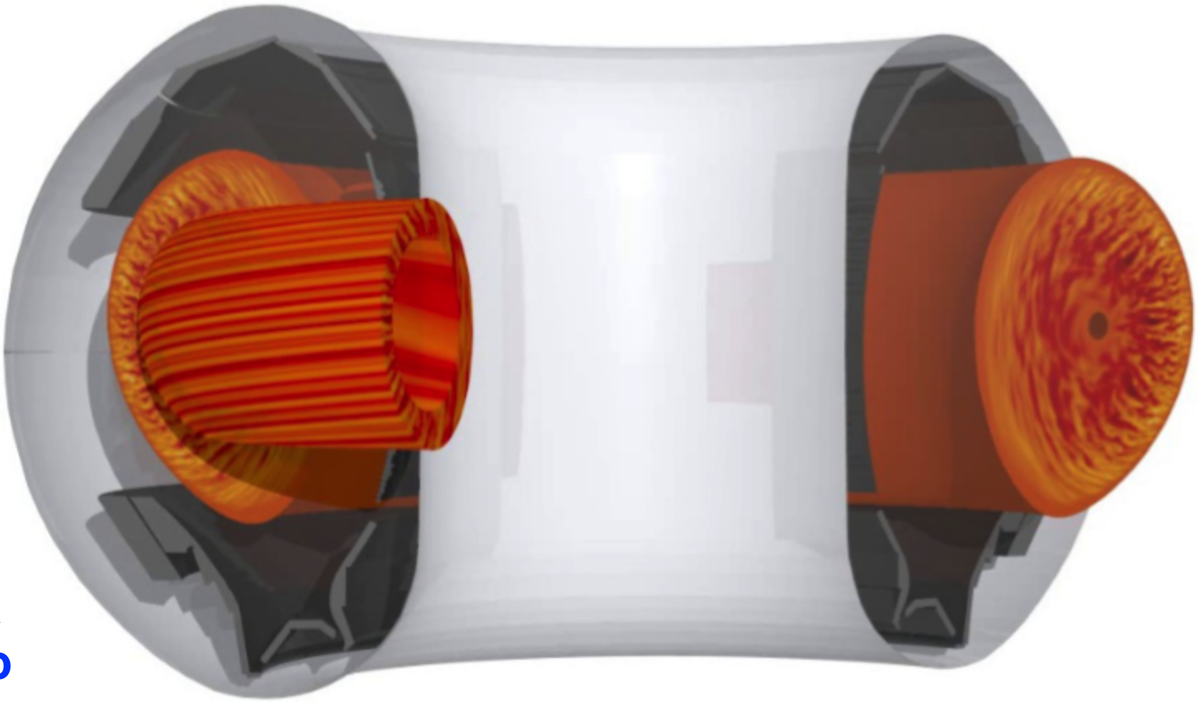

GENE gyrokinetic simulation of core turbulence

Frontier Multiscale Turbulent Transport Problems

Exploit Multiscale

- Bridging the gap between turbulence and transport time scales for global turbulence [TANGO code]
- Ion-electron multiscale turbulence [GENE, Gkeyll codes]
  - Practical algorithms for fast/efficient cross scale coupling [MuSHRooM code]
- Hermite-Laguerre gyrokinetic code [GX code]: seamless transition between fluid and kinetic to optimize rigor-efficiency

Deal Directly with Multiscale

- Turbulence in transport barriers (H-mode pedestal, Internal transport barriers) [GENE, Gkeyll codes]
- Ion-electron multiscale turbulence [GENE, Gkeyll codes]
  - Full multiscale simulations in H-mode pedestal
- New kinetic algorithms [Gkeyll]

What is the Gyrokinetic Equation?

- **Hyperbolic equation that describes time-evolution of the density of particles $f(x, y, z, v_{||}, \mu; t)$, in a 5D phase space – average over fast gyro-motion**
  - Plus parabolic diffusion-drag terms in a collision operator term *C[f]* & elliptic field solves
  - 3 spatial dimensions (x, y, z); 2 velocity coordinates $(v_{||}, \mu)$
  - Typical moderate-big grids: 256 x 256 x 32 x 16 x 8 (~300M points, 1000's of core-days)

- **Describes turbulence like 3D Navier-Stokes/MHD equations, but in a 5D phase space**

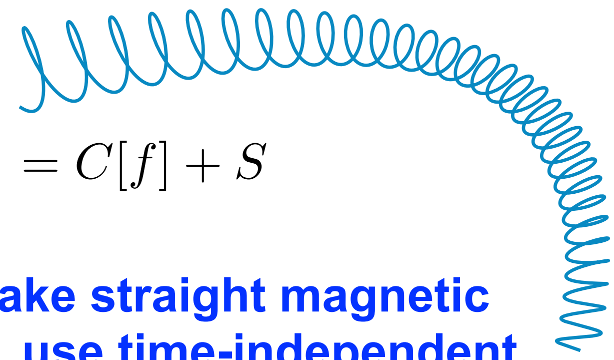

$$\frac{\partial f}{\partial t} + \frac{\partial}{\partial z}\left(v_{||} f\right) + \nabla \cdot \left(\vec{v}_E f\right) + \frac{\partial}{\partial v_{||}}\left(\frac{q}{m} E_{||} f\right) = C[f] + S$$

- **To illustrate, ignore gyro-radius (long wavelengths), take straight magnetic field, ignore magnetic fluctuations (pressure $\ll$ 8πB²), use time-independent dielectric coefficient $\epsilon_{\perp 0}(x)$:**

$$\vec{v}_E = \frac{1}{B^2}\vec{E} \times \vec{B}$$

$$\vec{E} = -\nabla \phi$$

$$-\nabla_\perp \cdot \left(\epsilon_{\perp 0} \nabla_\perp \phi\right) = 4\pi \sigma_{gc} = 4\pi \sum_s q \int d^3 v \, f$$

Poison Equation, charge q

$$= 4\pi \sum_s qB \int dv_{||} \int d\mu f_s(\vec{x}, v_{||}, \mu, t)$$

- **Can write the LHS of GK equation in Hamiltonian form**

G. Hammett

How do we solve the Gyrokinetic-Maxwell Equations?

$$\frac{\partial f}{\partial t} + \frac{\partial}{\partial z}\left(v_{||} f\right) + \nabla \cdot (\vec{v}_E f) + \frac{\partial}{\partial v_{||}}\left(\frac{q}{m} E_{||} f\right) = C[f] + S$$

- **Core Eulerian Gyrokinetic Codes: small 1% fluctuations f_1 about a Maxwellian equilibrium distribution f_0: $f = f_0 + f_1$**
- **Spatial gradients $\partial f_0 / \partial x$ in equilibrium distribution drive instabilities**
- **Turbulence is quasi-2D in (x, y) and extended along the magnetic field $\vec{B} = B\hat{z}$**

$$\vec{v}_E = \frac{\hat{z} \times \nabla\phi}{B} \qquad \nabla \cdot (\vec{v}_E f) = B^{-1}\{\phi, f_1\} + \vec{v}_E \cdot \nabla f_0 + \ldots$$

- **Nonlinear term can be written as Poissan Bracket** $\{\phi, \psi\} = \partial_x\phi\, \partial_y\psi - \partial_y\phi\, \partial_x\psi$
- **Typically use explicit or implicit time advance of $f\,(k_x, k_y, z, v_{||}, \mu; t)$**
- **Pseudo-Spectral nonlinear term {} evaluated in real space (2/3 de-aliasing rule)**
- **Lots of Bessel Functions J_0, J_1 for typical wavelengths $<$ gyro-radius**

Life at the edge is tougher:

- **Very steep gradients, atomic physics, $f_1 / f_0 \sim 30\%$: Global "full-f" for ion scales**
- **GKEYLL code is full-f, uses Discontinuous Galerkin methods**

Gyrokinetic Turbulence is Multiscale in 5 Dimensions

- **Kinetic turbulence: evolve f(x,v,t)**
  - Time scales: Turbulence vs Transport
  - Spatial scales:
    - Equilbrium vs fluctuations
    - Ion vs Electron
  - Phase space (scales in velocity space)
  - Low vs high order moments (i.e. fluid vs kinetic)
  - MGK addresses multiscale issues in all 5 dimensions

First Global Multiscale Simulation (GENE, TCV case)

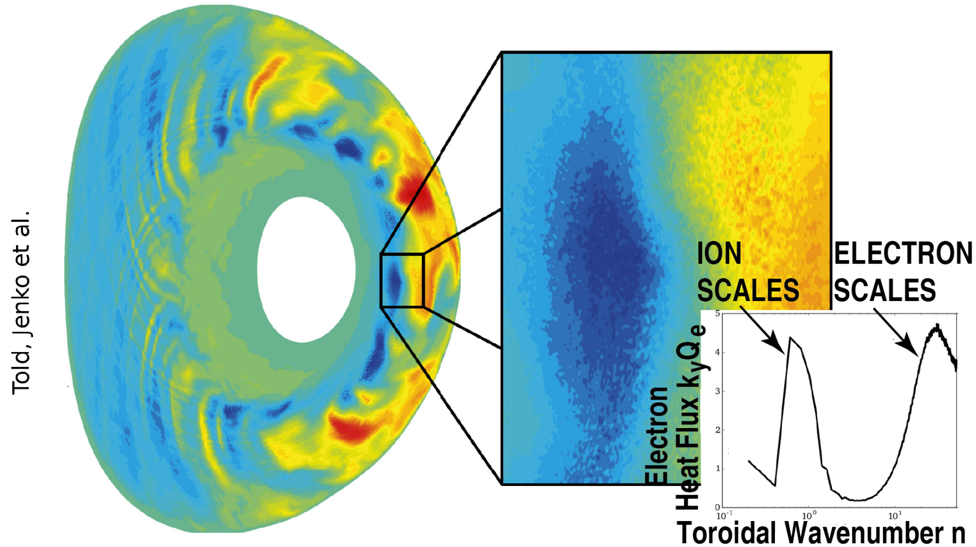

- **Multiscale phenomena present opportunity and challenge**
  - **Opportunity:** exploit scale separation to simplify
  - **Challenge:** need to address multiple scales directly

Practical Algorithms for Ion/Electron Multiscale Turbulence Interaction in Gyrokinetic Simulations

- **Exploit factor ~60 scale separation in y-direction and time to significantly reduce computation time for turbulent transport with cross-scale coupling**
  [Ernst, Francisquez, Pan, Hammett, Hatch, Dorland, Jenko]

- **Non-local coupling between scales important in some scenarios**
  [Maeyama, Howard et al. heroic ~30M hour gyrokinetic simulations]

Ion scales shear Electron scales

Electron scales damp ion scale zonal flows and beat into ion scales, increasing transport

- **Algorithm: Save time by not using fine electron y-grid & timestep for ion scales**

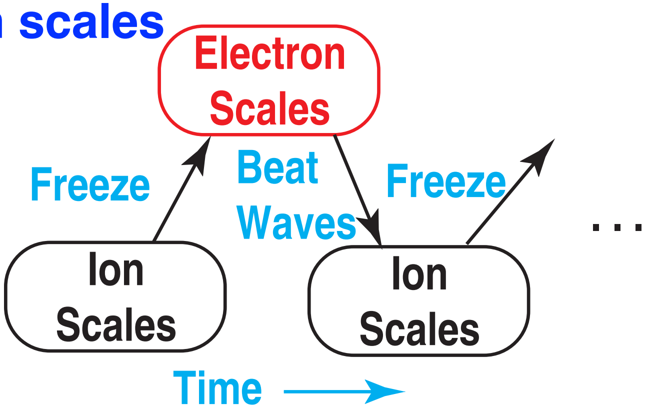

- **Use reduced model as test bed for time-stepping algorithms such as Additive Runge-Kutta [Reynolds]**

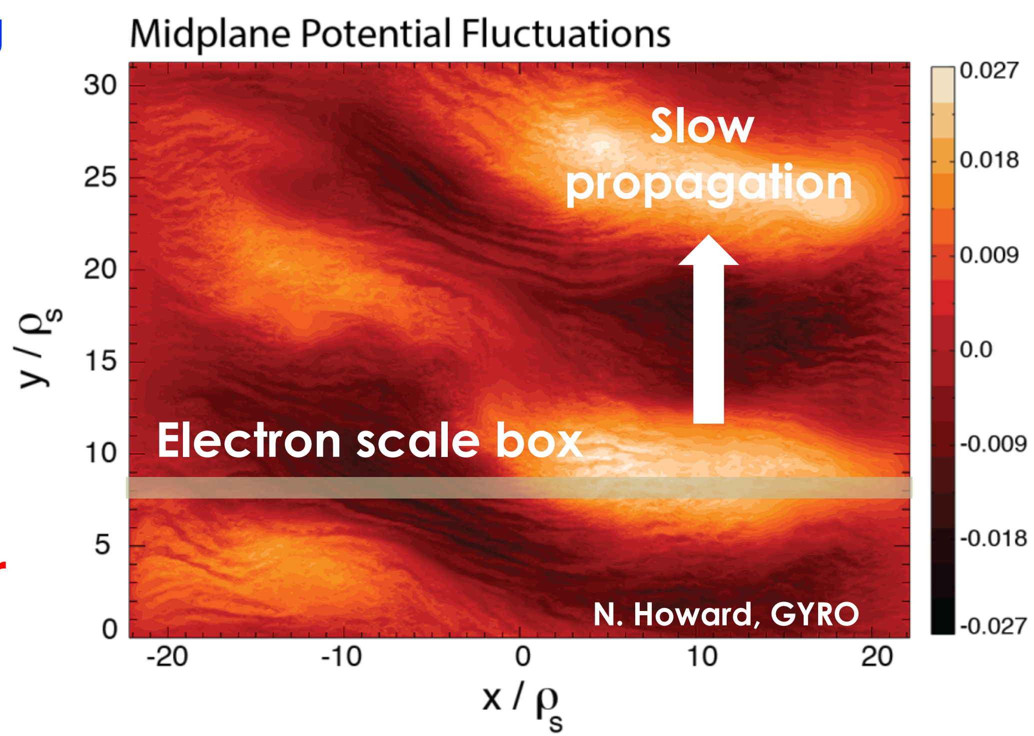

Formulated Reduced Model to Study Cross-Scale Coupling in 2D with 1-field

- **Nonlinear physics can be modeled with Modified Hasegawa-Mima Equation**

$$\frac{\partial \psi}{\partial t} + \{\varphi, \psi\} + \frac{\partial \varphi}{\partial y} = 0 \qquad \{\phi, \psi\} = \partial_x \phi \, \partial_y \psi - \partial_y \phi \, \partial_x \psi$$

$$\psi = \varphi - \langle \varphi \rangle_y - \delta_0(k) \partial_y \varphi - \nabla_\perp^2 \varphi$$

Growth inducing drive term (input)

- **Nonlinear term is Poisson Bracket of same form as gyrokinetic equation**

- ***Preliminary* model generalization to multiscale (2 orders of magnitude in k)**

$$\psi_{\mathbf{k}} = \left[\, \Gamma_0(k^2) + \Gamma_0(k^2\, m_e/m_i) - 2 + k^2 \lambda_{De}^2/\rho_s^2 - i k_y \delta_0/\rho_s \,\right] \varphi_{\mathbf{k}} - \Gamma_0(k^2)\, \varphi_{k_y=0}$$

$$\Gamma_0(k^2) = I_0(k^2)\ \mathrm{e}^{-k^2}$$

- **Bessel functions arise due finite gyroradius**

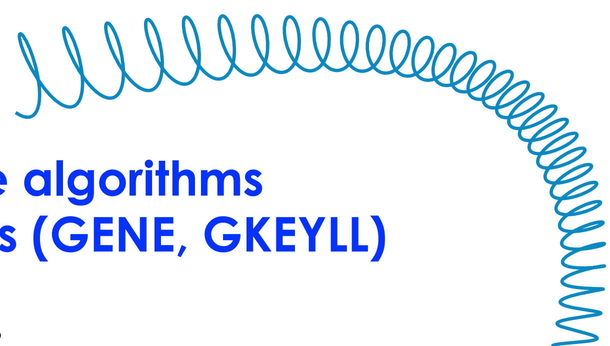

- **Expedient test-bed for multi-rate and multi-scale algorithms**
- **Transfer successful methods to gyrokinetic codes (GENE, GKEYLL)**

MuSHrooM Code to Test Scale-Separation Algorithms (with FASTMath ARKode library)

- **Standard 2D pseudo-spectral solver for $\varphi(t, k_x, k_y)$ using 2/3 de-aliasing rule, FFTW, and various time advance schemes, various hyperdiffusivities**

- **Use as test-bed for time-stepping algorithms**
- **Test ARKode for sub-cycling electron scales in collab. with D. Reynolds**
- **Expect dramatic speedups via higher order accuracy from ARK, allowing larger timesteps**
- **Opportunities for exploiting spatial scale separation in y**

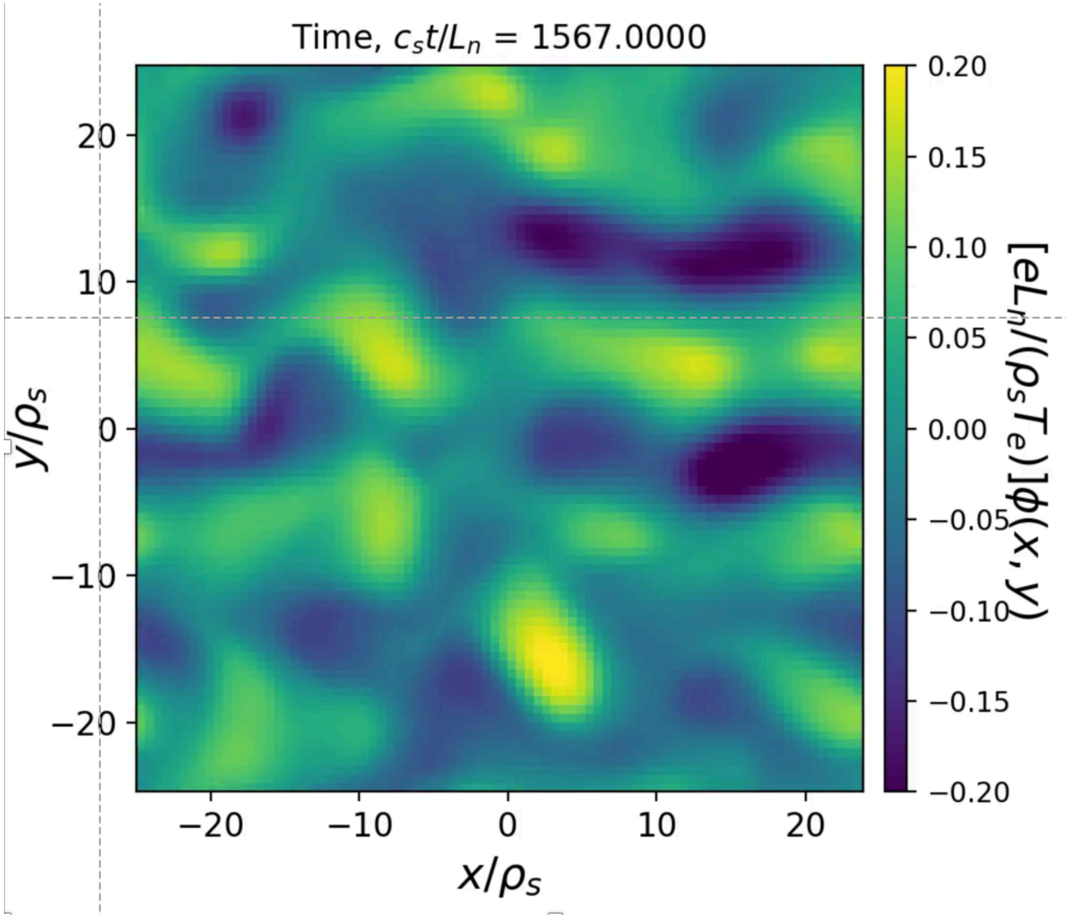

(Without ZF term)

M. Francisquez, Jay Lang, D. R. Ernst, G. W. Hammett

MuSHrooM Reproduces Earlier Ion-Scale Results

- **Same parameters and resolution used in [Smith & Hammett, Phys. Plasmas (1998)]**
  - Resolution is 128x128 in aliased real space, or 43x43 in k-space not counting negative k_x's

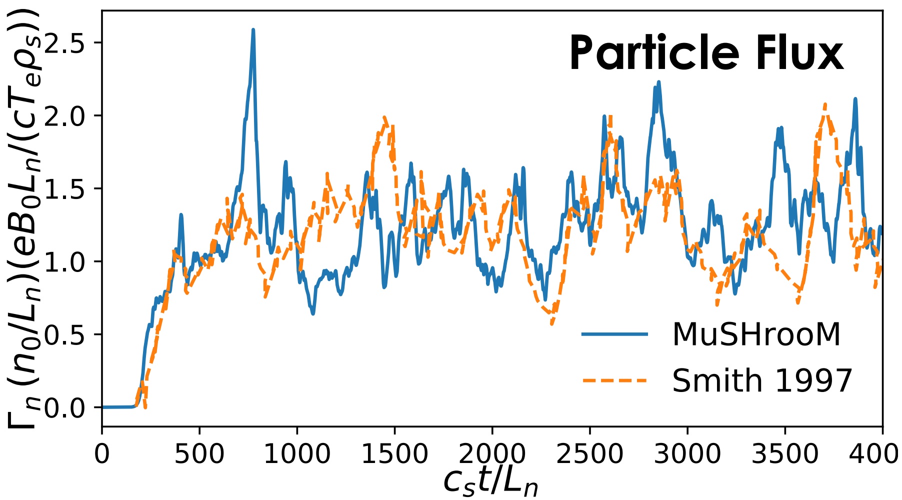

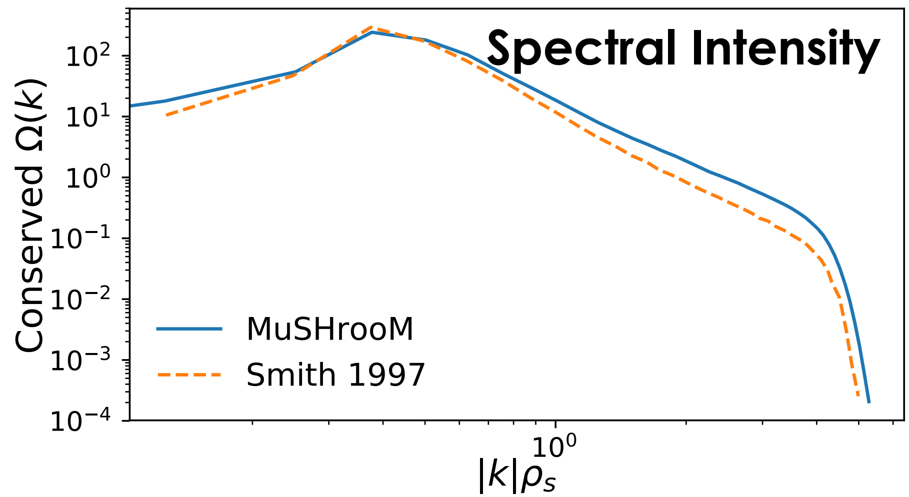

- **In both the spectrum and the growth rates the high-k trends are a little different.**
  - We do not have the same hyperdiffusion model.
  - Different time-stepping algorithm, which introduces a different form of diffusion.
  - Method for computing averages over k-space might differ.

M. Francisquez, Jay Lang, D. R. Ernst, G. W. Hammett

The Gyrokinetic GENE Code

Gyrokinetic Electromagnetic Numerical Experiment

- **Continuum approach to gyrokinetics (evolve distribution function on grid).**
- **Publicly available, world-wide user base from ~30 scientific institutions (US), ~100 worldwide**
- **Modes of operation:**
  - delta-f & full-f (gradient-driven, flux-driven)
  - flux-tube & full-flux-surface & global
- **Unique combination of various FDM and spectral methods**
- **Extensive physics: kinetic electrons, electromagnetic effects, collisions, realistic MHD equilibria, electron-scale turbulence…**
- **Part of fusion whole device modeling ECP project**

GENE on top-level HPC resources

PRACE

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

INCITE Award (2016)

Strong scaling of GENE on Titan (2k-16k nodes)

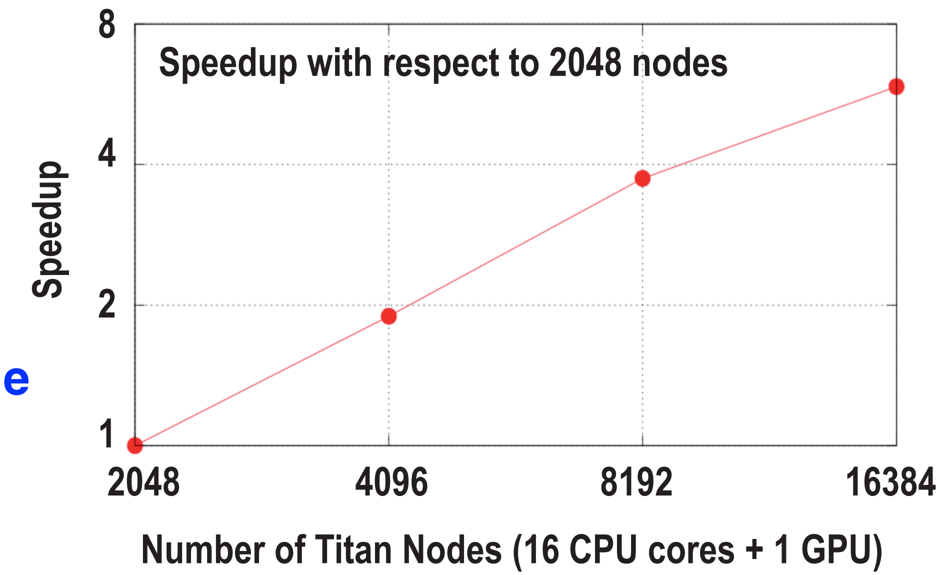

RAPIDS Collaboration for GENE Optimization

- **Working with Shan Hongzhang (LBNL) targeting Cori KNL**
- **GENE already scales well on Cori but improving on-node optimization**
- **Improved communication performance 15% by avoiding user defined MPI data types for non-contiguous data**
- **Extensive studies of FFT performance**
- **Work going forward will focus on GPU optimization**
  - RAPIDS work targeting NERSC 9
  - ECP targeting SUMMIT

- **Strong scaling of GENE on Cori**

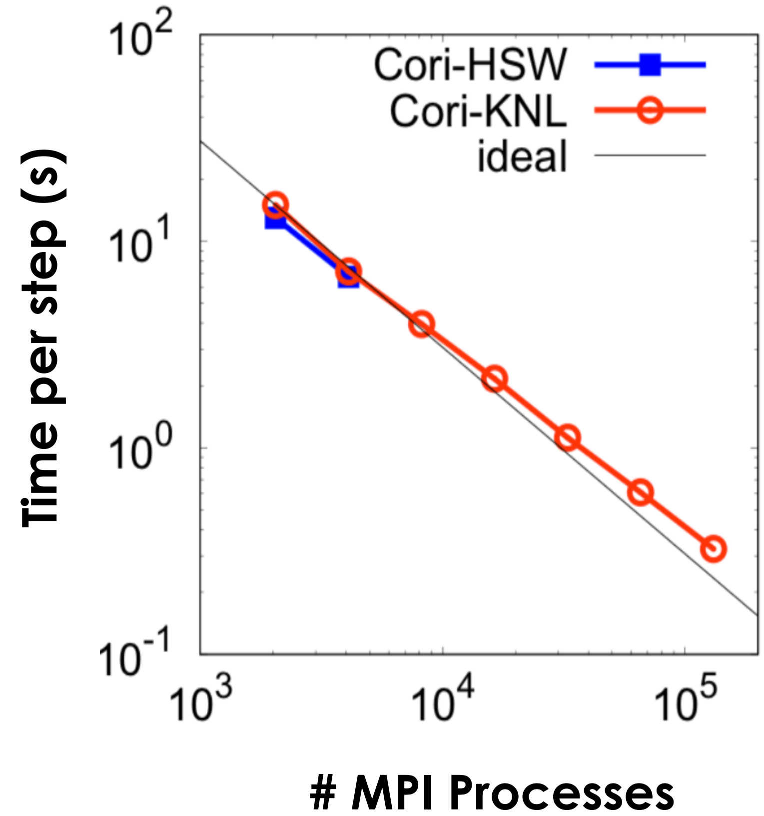

S. Hongzhang, F. Jenko, D. Hatch et al.

First Exact Gyrokinetic Linearized Landau Collision Operator in Conservative Form Formulated

- **Conservative and symmetric structure enables a finite-volume or spectral discretization that preserves the conservation laws**

$$C_{ab}^{\mathrm{gk}}\left(h_a, h_b\right)/\Gamma_{ab} = -\boldsymbol{\nabla}\cdot\left(\boldsymbol{J}_{ab}^T + \boldsymbol{J}_{ab}^F\right) + (\mathrm{FLR\ terms})$$

$$\boldsymbol{J}_{ab}^T = \int 2\pi d^2 v' \left(\frac{h_a}{m_b}\boldsymbol{I}_E^T\cdot\boldsymbol{\nabla}' f_{b0}' - \frac{f_{b0}'}{m_a}\boldsymbol{I}_D^T\cdot\boldsymbol{\nabla} h_a\right)$$

$$\boldsymbol{J}_{ab}^F = \int 2\pi d^2 v' \left(\frac{f_{a0}}{m_b}\boldsymbol{I}_E^F\cdot\boldsymbol{\nabla}' h_b' - \frac{h_b'}{m_a}\boldsymbol{I}_D^F\cdot\boldsymbol{\nabla} f_{a0}\right)$$

$$I_{\mu\nu} = \oint\frac{d\phi}{2\pi}\oint\frac{d\phi'}{2\pi} g_1\left(\phi\right) g_2\left(\phi'\right) g_3\left(k\rho'\sin\phi' - k\rho\sin\phi\right)\boldsymbol{e}_\mu\cdot\left(\frac{\boldsymbol{I}u^2 - \boldsymbol{u}\boldsymbol{u}}{u^3}\right)\cdot\boldsymbol{e}_\nu$$

$\boldsymbol{u} = \boldsymbol{v} - \boldsymbol{v}'$
$g_j\left(x\right) \in \{1, \sin x, \cos x\}$
$\mu,\nu \in \{\|, \perp, \phi, \|', \perp', \phi'\}$
$\Gamma_{ab} = 2\pi e_a^2 e_b^2 \ln\Lambda/m_a$

$\boldsymbol{I}_{E,D}^{T,F}$ **gyro-tensors for test-particle (field-particle) drag and diffusion coefficients (*can be pre-computed; constant in time*)**

Q. Pan and D. R. Ernst, Phys. Rev. E **99** (2019)

First Implementation of Gyrokinetic Exact Linearized Landau Collision Operator and Results

- Initially implemented[1] in the GENE gyrokinetic code using same finite-volume method[2] as the latest model operator (Sugama operator)
- Conservation to machine precision independent of resolution; H-theorem
- Zonal flow damping ~40% faster than Sugama model at high collisionality

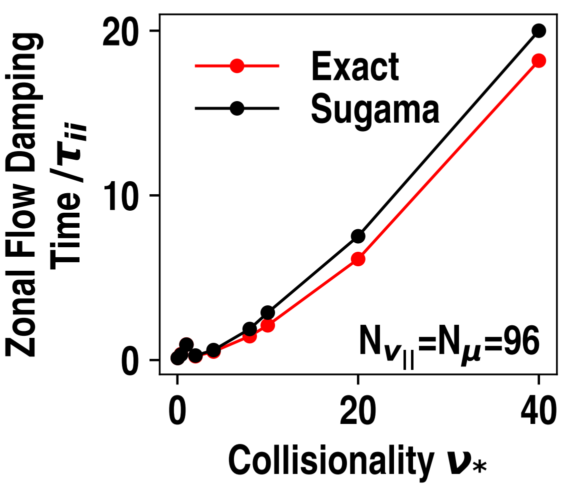

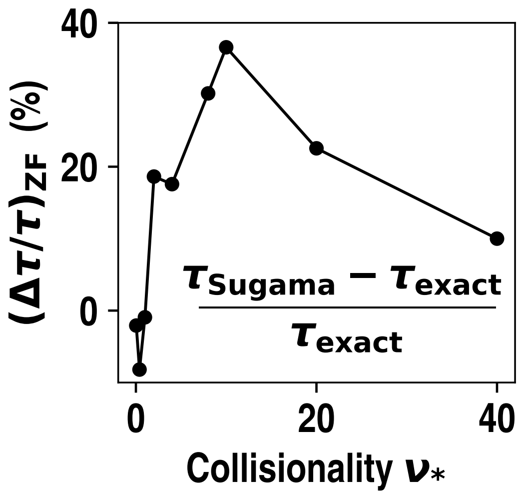

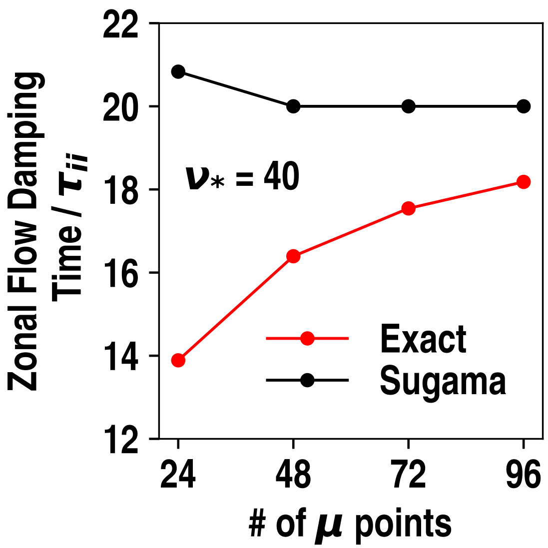

- Initial nonlinear runs completed comparing exact/model
- Opportunity to speed up with FASTMath & GPUs

Convergence with FLR effects requires >100 μ's (8-16 without)

[1]Q. Pan and D. R. Ernst, to be submitted
[2]P. Crandall et al., submitted to JCP (2018)

Potential Applications of FASTMATH Methods to Collisions

- **Spectral methods for 2D velocity space; significantly less grid points**
  - Formulate collision operator in weak (Galerkin) form

$$\frac{\partial}{\partial t}\int d^2v\phi_a h_a = \int d^2v \boldsymbol{\nabla}\phi_a \cdot \boldsymbol{J}_{ab} - \cancelto{0}{\int d^2v \boldsymbol{\nabla}\cdot\left(\phi_a \boldsymbol{J}_{ab}\right)}$$

  - Expand distributions and test functions in polynomial bases

$$h_s\left(\xi, v\right) = \sum_{l,m} \tilde{h}_{s,lm} P_l\left(\xi\right) T_m\left(v\right), \qquad \phi\left(\xi, v\right) = \sum_{l,m} \tilde{\phi}_{s,lm} P_l\left(\xi\right) T_m\left(v\right)$$

  - Use quadrature rule to perform velocity integration

$$\mathscr{M}\cdot\frac{\partial\tilde{\boldsymbol{h}}_a}{\partial t} = \mathscr{K}_{ab}^{T}\cdot\tilde{\boldsymbol{h}}_a + \mathscr{K}_{ab}^{F}\cdot\tilde{\boldsymbol{h}}_b$$

  - $\mathscr{M}$, $\mathscr{K}_{ab}^{T}$, $\mathscr{K}_{ab}^{F}$ are constant and precomputable; requires $\mathcal{O}\left(N^2\right)$ operations

- **Multi-rate time integrator for collisionless dynamics and collisions (tokamak core parameters); IMEX; Fast Multipole or Krylov-Implicit methods for collisions [D. Reynolds, C. Woodward]**

- **Nonlinear drift-kinetic implementation in GKEYLL edge code.**

Turbulence + Transport with TANGO: Iterative Solution

- **The Problem: solve for temperature profile $T(r)$ in a tokamak, balancing between input heating and turbulent heat losses**
  - Turbulent flux computed by GENE, which depends on $T(r)$, etc. (thousands of cores, tens of hours)
  - Balance equation is solved with a code TANGO

Nonlinear Problem

↓

Iterative Solution

Current solution method is a variant of fixed-point iteration

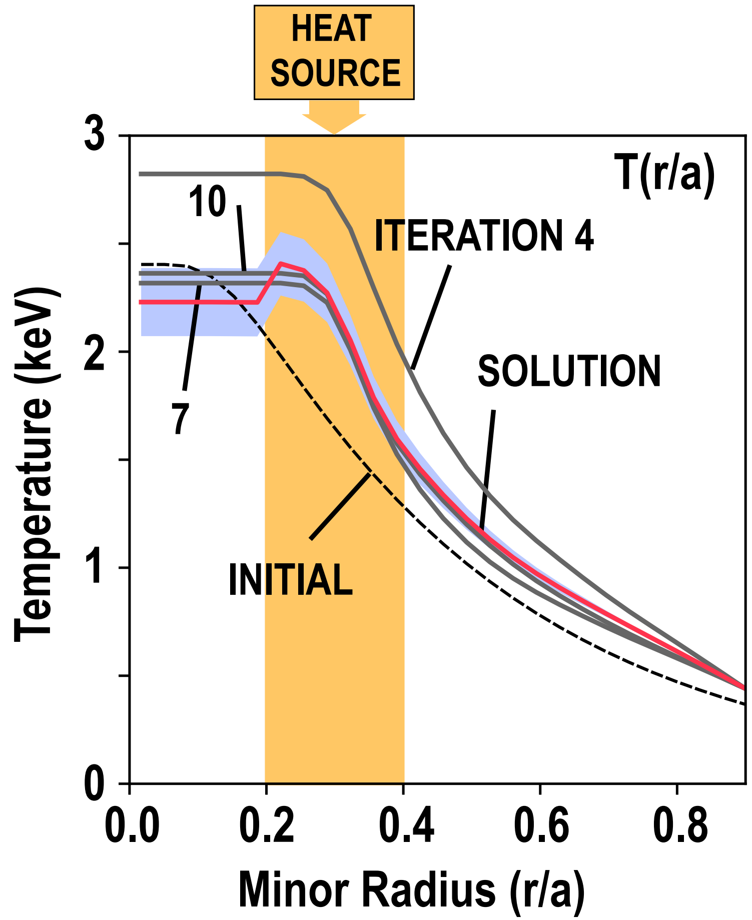

Parker, LoDestro, Told, Merlo, Ricketson, Campos, Jenko, Hittinger, *Nucl. Fusion*, 58 054004 (2018).

Jeff Parker (LLNL)
Lynda LoDestro (LLNL)

Potential FASTMath Collaboration: Advanced Iterative Methods in TANGO Transport Solver

- **Every iteration requires a call to a massive turbulence simulation**
  - 50-100 iterations in the course of a single GENE run are required currently (see Parker et al.)
  - Can be expensive!

- **Are there advanced methods, such as Anderson Acceleration, to reduce the number of iterations required?**

- **We've had discussions with Carol Woodward and David Gardner (LLNL) about using KINSOL's robust Anderson Acceleration solver**
  - Looks promising as an area to pursue

sundials

Jeff Parker (LLNL)
Lynda LoDestro (LLNL)

Turbulence in Edge Pedestal

- **Burning plasmas will need an edge transport barrier**
  - ITER, e.g., will need ~4 keV pedestal
  - Need to understand and predict
  - Recent breakthroughs in pedestal turbulence with GENE
    [Hatch et al NF 2016, 2017, 2018, 2019; Kotschenreuther et al NF 2017, 2019]

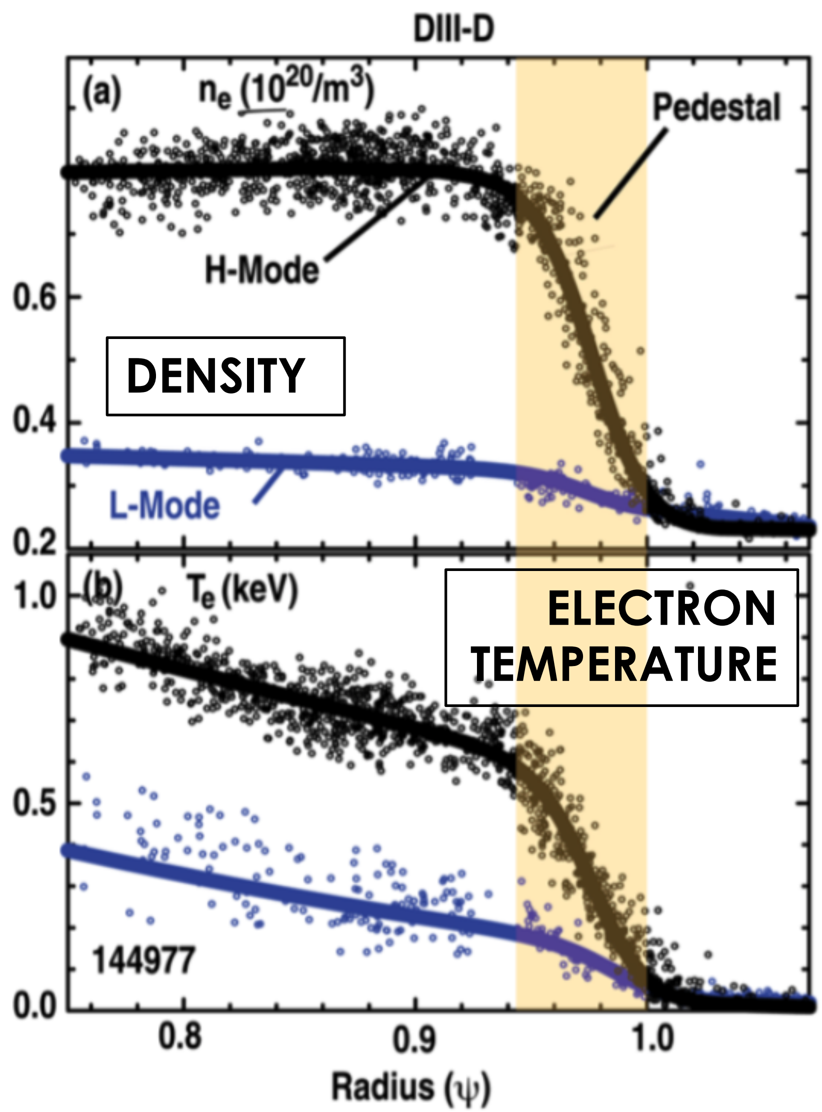

- **Goals**
  - Develop reduced models / surrogates
  - UQ with reduced models
  - Ion-Electron multiscale treatments

D. Hatch, M. Kotschenreuther (Univ. Texas)

Ion-Electron Multiscale Turbulence in the Pedestal

Reliable reduced models for pedestal transport currently do not exist

- **Map linear simulations (cheap) to nonlinear transport levels (train on high-fidelity simulation data)**
- **Further refine in comparison with experimental data**
- **Needs:**
  - Machine learning or advanced statistical techniques to define mapping
  - Efficient algorithms for exploring a high-dimensional parameter space
  - ML / statistical approaches to learning from both high-fidelity sims and experimental observations

- **Ion-electron multiscale: substantial transport at ion scales and electron scales.**
- **So far no tests of cross-scale coupling (likely very different than core multiscale).**

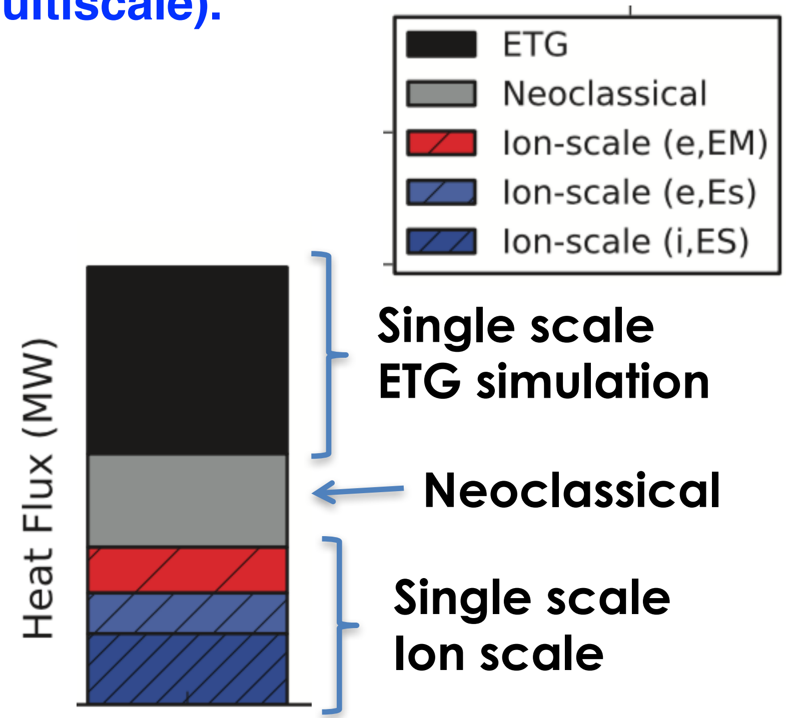

D. Hatch, M. Kotschenreuther (Univ. Texas)

Gkeyll: First Continuum 5D Gyrokinetic Simulations of Turbulence in SOL with sheath model boundary conditions

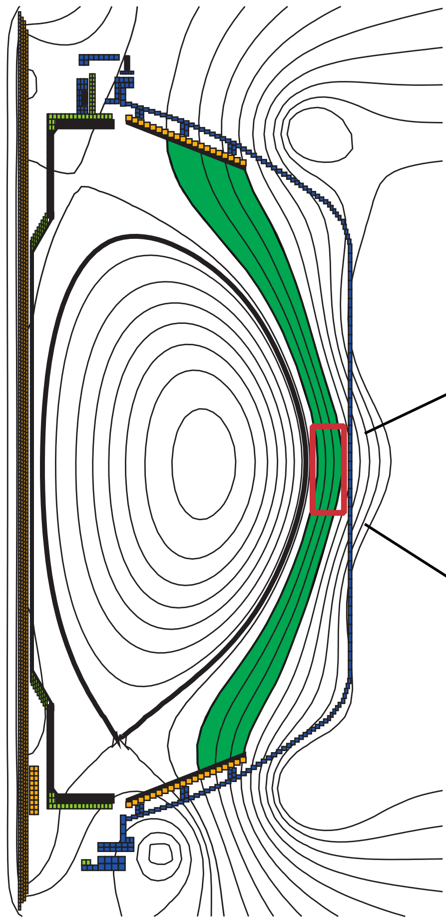

Edge region very difficult computationally.

XGC only code crossing last closed flux surface at present.

MGK SciDAC
G. W. Hammett,
A. Hakim,
M. Francisquez,
et al.

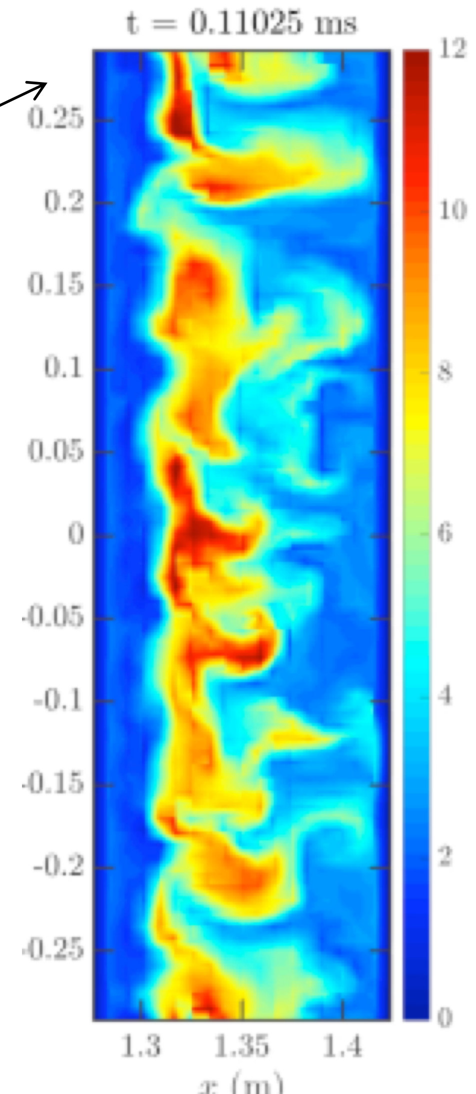

Simplified helical model of SOL (toroidal + vertical B field) – general B-field geometry under test.

E. Shi, A. Hakim, G. Hammett, et al. J. Plasma Physics (2017), T. Bernard et al. Phys. Plasmas (2019).
Q. Pan et al. Phys. Plasmas (2018), similar work in straight fields with GENE.

Gkeyll Code: Novel Kinetic Algorithms, Multiple SciDACS

- **Novel version of Discontinuous Galerkin (DG) algorithm**
  - Conserves energy for Hamiltonian system even with upwind fluxes [Juno, Hakim, et al. JCP 2018]
  - High-order local algorithms reduce communication costs, helpful for Exascale.
- **New modal version ~30x faster than nodal version**
  - Computer-algebra generated code (w/ Maxima) makes use of sparseness of modal interactions.
- **Framework: LuaJIT over C++, uses ADIOS, Eigen, MPI, …**
- **3 Main Versions, used in 3 SciDACs:**

1. **Gyrokinetic DG version** for edge turbulence in fusion,
   in MGK SciDAC project (D. Hatch, PI) for pedestal / multiscale work,
   in HBPS SciDAC project (C.S. Chang, PI) for scrape-off-layer turbulence work.
2. **Vlasov-Maxwell/Poisson DG version:** solar wind turbulence (PU/Maryland), plasma-surface interactions in thrusters (AFOSR / Virginia Tech) & tokamak disruption SciDAC (LANL / Virginia Tech)
3. **Multi-moment multi-fluid (extended MHD) finite-volume version:** reconnection (Princeton Center for Heliophysics), global magnetosphere simulations (UNH)

Gkeyll Code: opportunities for FastMath / Computer Science / Applied Math Collaborations

- **In HBPS SciDAC, collaborating w. ORNL applied math group (Cory Hauck & Eirik Endeve) on RK-Legendre super-time stepping for collision operator (diffusion-advection operator in velocity coordinates).**

Many opportunites for more collaborations:

- **Efficient implicit solver for parallel electron dynamics.**
  - What would work? Most iterative implicit methods for parabolic/elliptic problems, but fast electron motion is a hyperbolic operator: high-frequency wave needs to be implicit for stability, but not important itself.
- **Improve scalability of IO—issues beyond a few thousand cores. Want to discuss with ADIOS group.**
- **Help with performance tuning, starting with hotspot analysis.**
- **Porting to GPU**
- **Post-processing, visualization, in-situ analysis**
- **Better parallel Poisson solvers when we get to larger systems (was using PETSc, now Eigen)**

GX code: Seamless Transition from Fluid to Kinetic

- **GX code (Dorland et al)**
  - Laguerre / Hermite in velocity space
  - Can use intelligent closures and be gyrofluid when applicable
  - Alternatively can keep lots of moments and be fully gyrokinetic
  - Status:
    - Collision operator implemented
    - Linear closures tested
    - Extensive linear benchmarks
    - Nonlinear version runs on GPUs
- **Developing ETG edge model**
- **Spectral deferred correction in collaboration with M. Minion (LBNL)**

[Mandell, Dorland, Landreman, JPP 2018]

- **Goals:**
  - Seamless transition between fluid and kinetic descriptions
    - Depending on physical regime, desired speed / accuracy
  - Use machine learning to train fluid closures from kinetic simulations
  - Start with toy problem then apply to GX

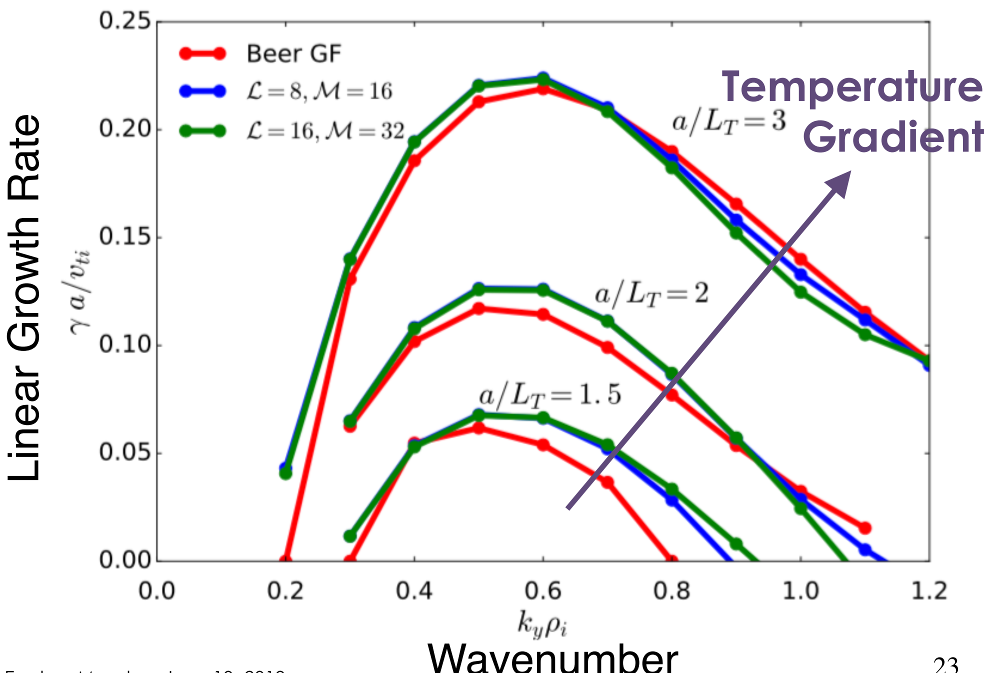

Fusion MGK Database

Unified and cross-referenced database of experimental and model sensitivities

- **Cross-validation studies between**
  - Simulation codes
  - Simulations and experiments
  - Different experiments
- **Advanced data analytics**
  - Real-time diagnostic tools
  - Machine learning algorithms
  - Data mining of simulations and experiments
- **Consistent and rigorous uncertainty quantification**
- **Multifidelity hierarchy design**
- **Model adequacy and prediction assessment capabilities**
- **Validation, optimization, design, data mining**

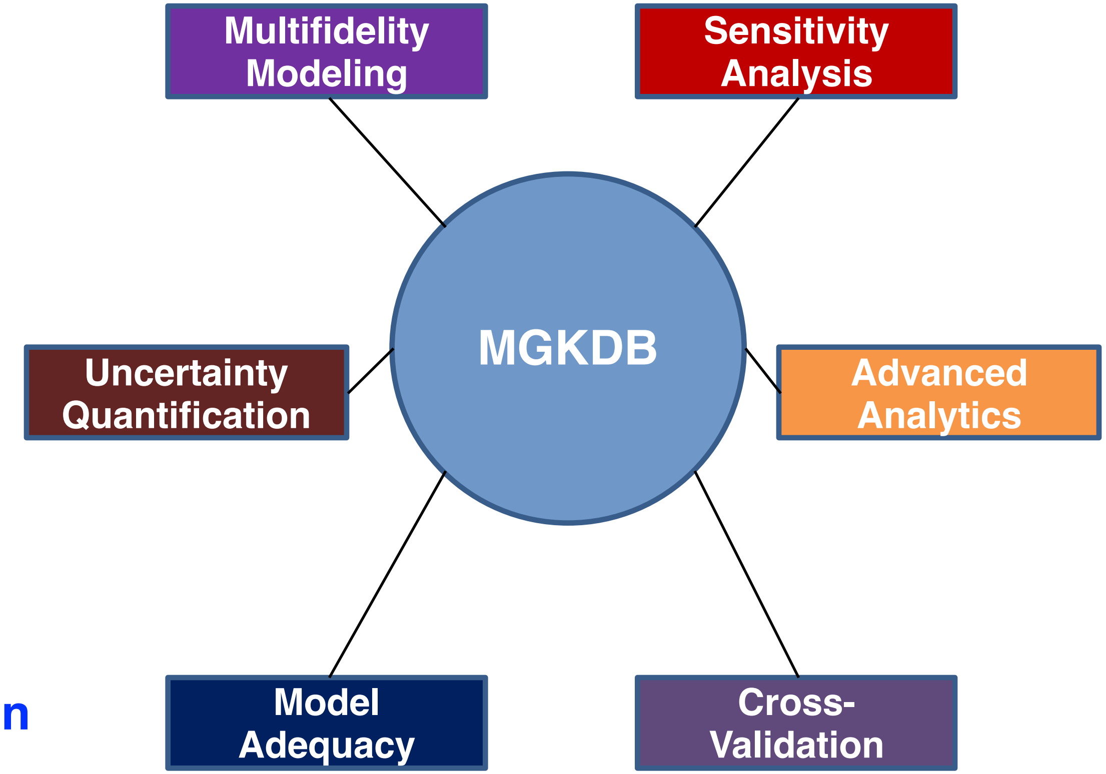

Craig Michoski (ICES, Univ. Texas)

Current Status of MGKDB, RAPIDS connection

- **Postdoc on the way**
- **Early stages. MongoDB prototype developed. Running locally and on NERSC.**
- **RAPIDS (i.e. Tahsin Kurc, Kshitij Mehta, …) interested in interfacing a DB through ADIOS for performance and large data transfer, and helping to facilitate this**
- **Existing codes already connected with MGKDB (e.g. GENE, GS2, GX, GYRO, QuallKiz, etc.)**
- **Existing Open Source European initiative from Jonathon Citrin, Karel van de Plassche, Yann Camenen, called GKDB:** https://github.com/gkdb/gkdb.git
  - Focused on ITER Integrated Modelling & Analysis Suite (IMAS)
  - OMFIT integration
  - PostgreSQL relational database
  - Strong interest in combining forces to maximize utility/workforce for the community at large

Craig Michoski (ICES, Univ. Texas)

Summary of FASTMath and RAPIDS Collaborations

- **[MuSHrooM, GENE] D. Reynolds (SMU), C. Woodward (LLNL)**
  - Multi-rate methods (ARKode) for multiscale turbulence, collisions
- **[TANGO] Carol Woodward and David Gardner (LLNL)**
  - KINSOL's Anderson Acceleration solver
- **[GX] M. Minion (LBNL)**
  - Spectral Deferred Correction in fluid-kinetic code
- **[GENE] Shan Hongzhang (LBNL)**
  - Optimizing Cori KNL
- **[MGKDB] Tahsin Kurc, Kshitij Mehta, … (RAPIDS)**
  - Interface DB through ADIOS for performance and large data transfer

Opportunities for Future Collaboration (i.e., Needs)

- Spatial scale separation in multi-scale turbulence
- Implicit Krylov methods for collisions
- Machine learning or advanced statistical techniques to define mapping
- Efficient algorithms for exploring a high-dimensional parameter space
- Efficient implicit solver for parallel electron dynamics.
- Improve scalability of IO [ADIOS]
- Help with performance tuning, starting with hotspot analysis.
- Porting to GPU
- Post-processing, visualization, in-situ analysis
- Better parallel Poisson solvers

Thank You !